Microsoft Metadirectory Services Concepts and Architecture

Operating System

Abstract

This document provides an overview of the capabilities and concepts behind Microsoft? Metadirectory S relationship to the concept of identity management.

The Problem of Identity Management

The metadirectory provides a solution to the problem of identity management.

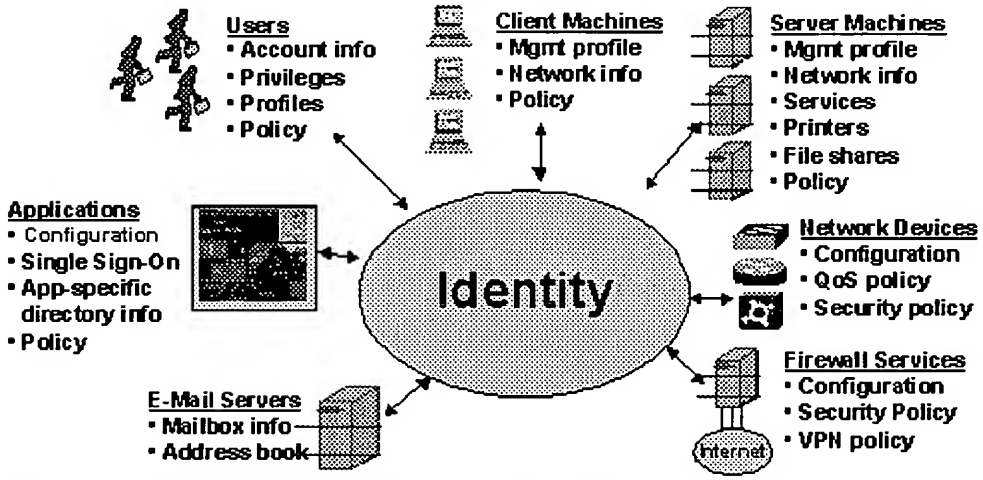

Figure 1 The Identity Management Challenge

As illustrated in Figure 1, identity is the summary of information about people, applications, and resourc directories and databases throughout most IT enterprises. Examples of identity data associated with peop mailboxes, salaries and job titles. Application identity information includes the network addresses where servers. It also includes lists of services that applications can provide. Network resources, such as printer attributes ? their location and the printing capabilities they support, for example.

The Identity Management Challenge

The diversity of identity data and the number of places where such data reside raise a number of manage

- Not all identity data is kept in directories or exposed through a directory service interface such as l Directory Access Protocol (LDAP). For example, many systems only expose identity information application programming interfaces (APIs).
- Identity information frequently is duplicated in multiple places, and versions tend to drift out of sy time if left unchecked.
- Typically, there is no single place where administrators and applications can access or manage an . (sometimes called a join) of an enterprise's identity information.

- The number of places where companies must manage identity data increases with each additional platform.

These challenges make it difficult for companies to implement comprehensive and integrated identity ma Leaving an enterprise environment in this state increases cost and complexity.

Common Identity Management Scenarios

Most large companies are already starting to grapple with some form of identity management project. Cc include:

- **Global address book applications.** Synchronizing mailbox information between the different e-m a company enables users to locate other users and send them e-mail across differing systems.
- **Hire/fire solutions.** Propagating information about a newly hired employee ? such as title, role an all systems that require identity data enables speedy establishment of services. Systems also must processes quickly in reverse when employees leave to prevent breaches of security.
- **E-commerce applications.** Synchronizing enterprise identity information, such as digital certifica extranet users, is enabled with directories that reside outside of firewalls.

Solution Requirements

In the past, many companies have tried to create a single directory to hold all enterprise identity informa efforts failed for several simple reasons:

- Many applications cannot be modified easily to use directories.
- There are good reasons, such as various replication and security requirements, why some applicati identity in their own formats.
- Political boundaries inhibit complete consolidation regardless of what is technically possible.

This suggests that identity data will continue to exist in many places, and companies need to find ways t directory services and application repositories work together. Assuming that there will be many identity solutions must provide:

- Connectivity to many forms of identity data.
- Management of identity flow between repositories.
- Mechanisms for maintaining data integrity throughout the identity management infrastructure.

We discuss each of these issues in more detail below.

Connectivity

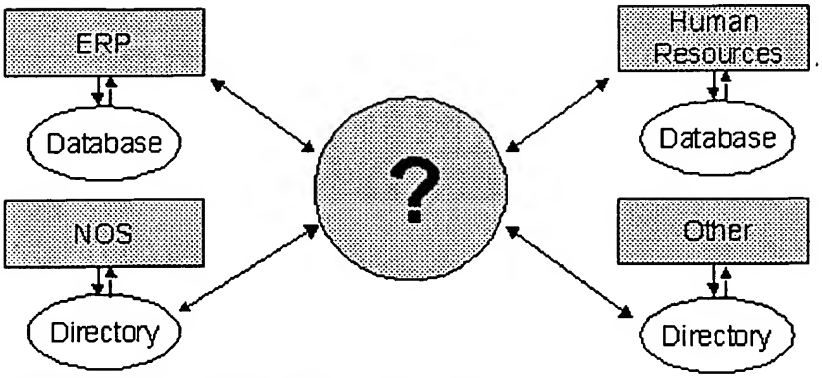

Figure 2 Connectivity Requirements

Connectivity requirements are simple: the more directory services, databases and applications to which i solutions can connect, the more value they can offer. As illustrated in Figure 2 above, unknown data in o obtained from another. An identity management solution can connect to a given repository, if it is able tc

- Obtain information about what has changed in the repository.
- Add new objects to the repository.
- Delete objects from the repository.
- Change an existing object's attributes to different values.

To be a comprehensive solution, technologies should be able to connect to data in:

- Standards-based directory services via LDAP Version 3.
- Popular existing e-mail applications and non-LDAP directory services.
- Enterprise Resource Planning (ERP) applications.
- Databases via access methods such as SQL.
- Applications in which the only interface to identity information is through application programmir and no directory interface is available.

Information Management Flow

Information management flow is the process of managing the flow of identity information between repo: management flow functionality must be able to:

- Detect changes to identity data and propagate updates to other repositories.
- Aggregate data from different repositories into metadirectories that contain a holistic view of ident the enterprise.

Track related objects as they change their positions in directory trees and other repositories due to period

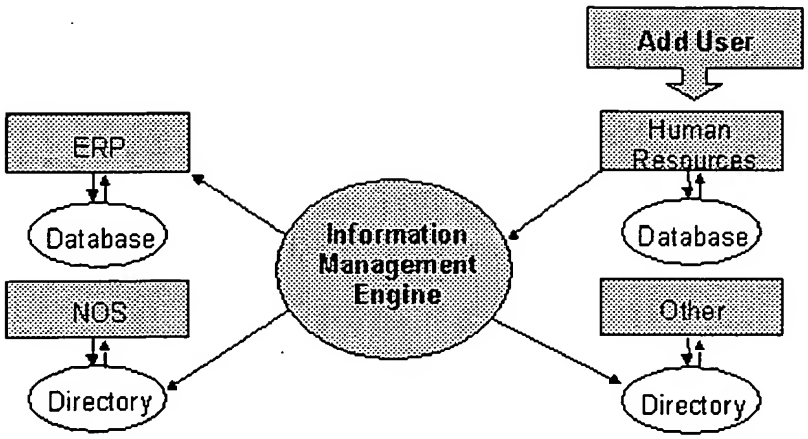

Figure 3 Change Event Processing

Change Event Processing

Change events occur any time administrators, users or applications add, delete or modify a piece of ident repository. Unmanaged, identity data changes quickly becomes disorganized. Identity management solut provide features to detect changes, perform necessary data format translations and then update all reposit reflect the changes. For example, if an administrator adds a new employee to the human resources (HR) event needs to cause systems that the person will use to reflect the addition. In Figure 3 above, the chang other directories and applications.

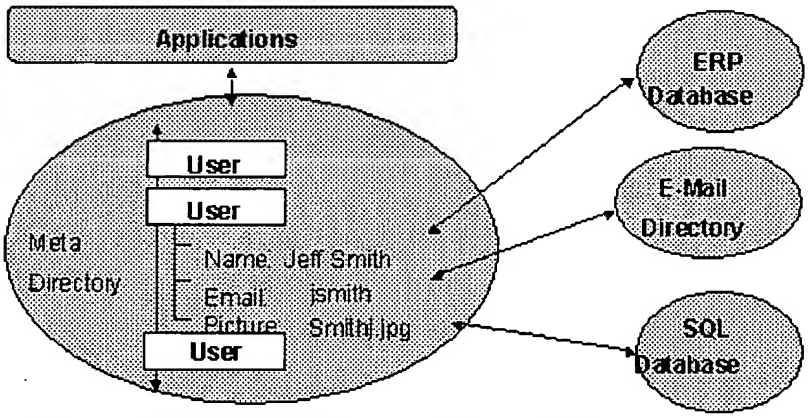

Figure 4 Data Aggregation in a Metadirectory

Data Aggregation Capabilities

While identity information resides throughout most enterprises, directories that contain an aggregation o many other repositories can offer great value. This metadirectory concept was pioneered by The Burton the term *join* to represent an aggregated view of an enterprise's identity data.

With a metadirectory, applications can access a variety of information in one place, using a single access model, instead of interacting with each of the source repositories. Metadirectories also maximize perforn can be stored in indexed form. There is no need to fetch data from sources, which may reside across wid (WAN) connections, at runtime. To offer the greatest value, data aggregation capabilities must be able tc

- Gather and incorporate information from many sources including directories, databases and applic

- Group related information together even though it may be stored in different ways in different plac about a user named Jeff Smith might be stored under names such as Jeff Smith, jsmith and smithj i as seen in Figure 4 above.
- Push changes back out to sources when users or applications make changes to the aggregated view metadirectories must be integrated with change event processing infrastructures.

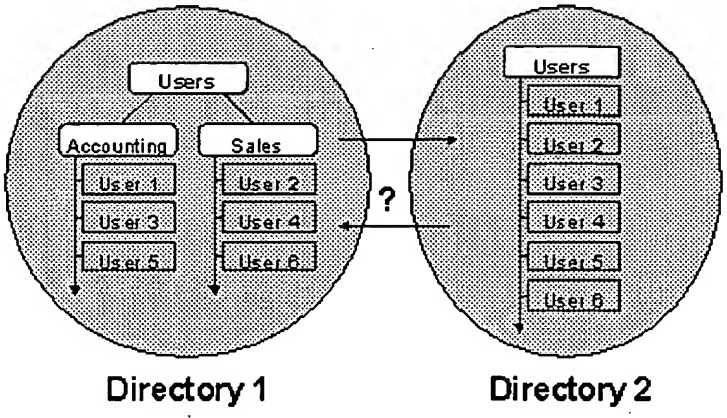

Figure 5 Tracking Related Objects

Related Object Tracking

When administrators deploy identity management solutions, they must be able to tell the identity manag that Jeff Smith, jsmith and smithj are all the same person. Then, as seen in Figure 5, the engine must be a relationships as identity data is periodically reorganized. Solutions must not lose track of users simply be position in a directory tree structure moving from the Accounting department to the Sales group, for exa

Integrity Management

Integrity management is the process of ensuring that identity data does not become corrupt or out of sync repositories as changes occur. Integrity management functionality must be able to:

- Maintain identity data ownership relationships.
- Act appropriately when failures occur.
- Maintain referential integrity within identity data.

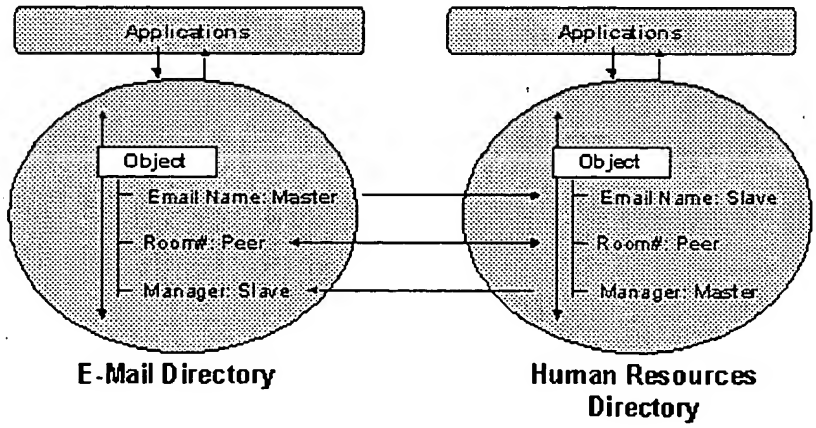

Figure 6 Managing Ownership Relationships

Ownership

An important aspect of enterprise identity management is recognizing ownership relationships that must between applications and data. For example, a person's mailbox name is owned by the e-mail system tha Within most companies, the HR system owns the data corresponding to whether or not a person is an act no enterprise identity management infrastructure in place, these ownership relationships are preserved by other applications have the ability to access and update e-mail and HR data. With synchronization conne flow management deployed, however, the situation changes.

- Consider a case in which mailbox information is being synchronized with the HR directory by a c in Figure 6 above. If the connector is not configured correctly, a user could change the mailbox att system and the connector would overwrite the mailbox value in the e-mail directory, causing trem Solving the problem is not as simple as just preventing changes from flowing backwards to the e-r HR system may own information, such as the name of a person's manager, which must flow back t directory. Other attributes, such as a person's office number, may have no clearly defined ownersh be data that anyone can update.

As a solution requirement, administrators must be able to define and enforce ownership relationships at t change is in accordance with the ownership rules, it is allowed to pass through, otherwise it is blocked o example, if a person changed a mailbox attribute in the HR directory, the identity management solution v attribute back to the value contained in the e-mail directory.

Failure Management

The ability to propagate a change to multiple repositories is a key requirement for identity flow managen Yet, any time an engine makes multiple updates, the opportunity exists for one or more of the updates to different repositories to become inconsistent as illustrated in Figure 7 below. For example, if a person's t spending limit are changed, but the metadirectory is unable to update the user's title in applications, iden in a state of confusion. Typically, this means that an administrator must investigate the situation and mak

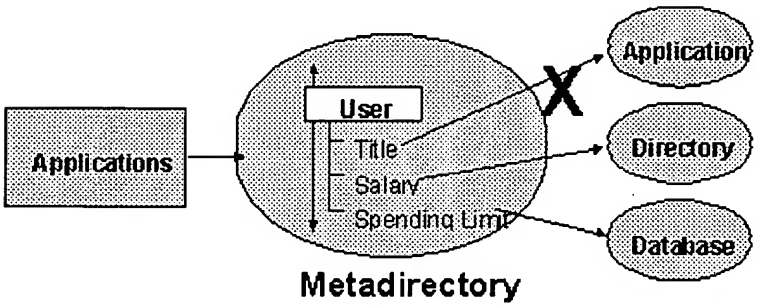

Figure 7 Managing Failures and Maintaining Referential Integrity

In database systems, this challenge is usually addressed with transactions that ensure all updates occur su rolled back as a unit. Unfortunately, most directory services and application programming interfaces do n transactions. This means that identity management solutions must find other ways ? such as using log-ba mechanisms that continue to request changes until confirmed ? to ensure that all repositories eventually n

Referential Integrity

Another challenge that identity management solutions share with databases is maintaining referential int repositories. Referential integrity refers to the need to maintain relationships between the values of relate different locations. For example, identity management solutions must be able to ensure that a person's tit resources system is consistent with the person's spending limit in the procurement system. Databases sol providing stored procedure and trigger features that enable administrators to execute a business rule eacl changes. Directory services do not provide similar features today. Therefore, identity management soluti capability to execute business rules, which will reject changes that do not meet referential integrity requi

Only a metadirectory solution addresses all these issues.

The Metadirectory Solution

If Internet/intranet, proprietary e-mail, and other directories contain identity information about only *som* the metadirectory is capable of containing identity information about *everybody everywhere*. The metadi integrate any number of disparate identity repositories in virtually any format. Thus, the metadirectory b root of identity information within the enterprise. The metadirectory provides the rationalized and unifie objects that consist of attributes from a variety of directories. This integration enables you to lower admi eliminate duplication, reduce discrepancies, and make the identity information widely available. The me enough to adapt itself to any enterprise's organization, structure, politics, and management styles; and dy change as they change.

Sources

The metadirectory collects its identity information from the other connected directories and repositories Nearly all e-mail, database, and other directory applications can export their contents in some form. The collect this data through file exchange, in an e-mail message, or through an on-line, protocol-driven tran administrator or end user can add other metadirectory identity information.

Content

We usually think of directories as containing identity information about people, such as e-mail addresses view. The metadirectory can contain much more information about any real-world objects. Objects may

- Physical, such as people or computers;
- Conceptual, such as organizations or departments;
- Geographic, such as countries or cities;
- Digital, such as document files for on-line viewing.

The only requirement of the metadirectory is that these objects be organized in some sort of hierarchical example, a person might be described as part of a department that is part of an organization that is locate domain or a country. Or, in a multi-national corporation, an employee might be part of a division located falls under the corporation in the organizational tree.

A person is not necessarily the lowest level of the hierarchy. For example, a document or a portable com that person might also be represented by a directory entry below the person entry in the tree.

Management

The management of metadirectory contents and security can be centralized, distributed, or a combinatior metadirectory can be created so that changes to certain entries can be made only in the connected directo the metadirectory. Changes to other entries may be made only in the metadirectory and then propagated directory. Different people can manage different portions of the metadirectory. This level of control exte entries themselves, but also to the individual attributes. Therefore, end users can manage parts of their o information ? telephone numbers or addresses, for example. The metadirectory does not impose any mar lets you create a directory whose management matches the realities of your organization, its security and requirements.

Microsoft's Metadirectory

Microsoft has a metadirectory solution that has already been widely used to meet the challenges of enter management.

Its Origins

In July 1999, Microsoft purchased ZOOMIT Corporation. The ZOOMIT Corporation was known as the delivering a metadirectory solution. Through this purchase, Microsoft is able to provide a comprehensive platform for Microsoft? Windows? 2000 Server that includes Windows security, the Active Directory? s metadirectory services. The ZOOMIT metadirectory solution addresses the problems discussed earlier in time, Microsoft's metadirectory solution, Microsoft Metadirectory Services, will be completely integrate distributed systems offerings, making it an even more powerful identity management tool.

Its History

Microsoft Metadirectory Services is thus an established product with a long history and an extensive and base. The ZOOMIT Corporation began shipping their ZOOMIT VIA 1.0 product in October of 1996.

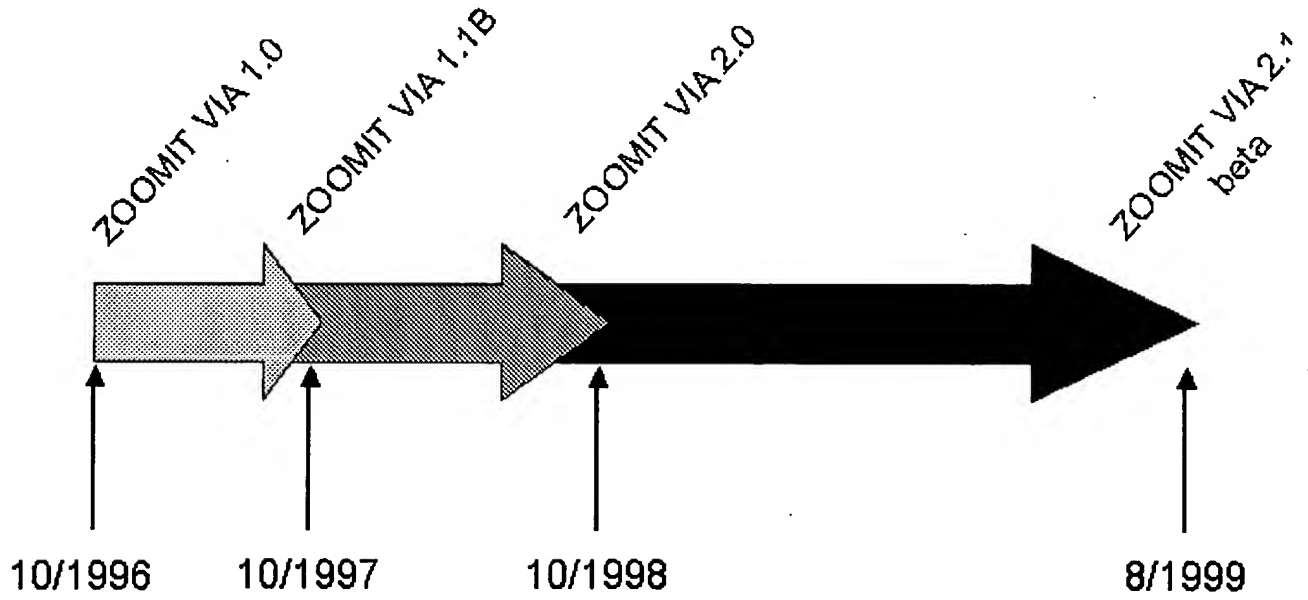

Figure 8 ZOOMIT VIA Development Timeline

A beta version of ZOOMIT VIA 2.1 was shipped to several customers during the transition to Microsoft

release and subsequent releases will be referred to as Microsoft Metadirectory Services (MMS). Since M ZOOMIT, two additional versions of MMS have been released: MMS version 2.1 in December 1999 anc 2000. Many large organizations throughout the world now successfully use MMS in complex and demar

The remainder of this document focuses on the current version, MMS 2.2. It covers the basic concepts of a flexible and powerful architecture, and shows how it can be used to solve complex, real-world problem management.

Microsoft Metadirectory Services

MMS provides an industry leading solution for the identity management problems such as enterprise adc hire/fire scenarios. Conceptually, the components of the metadirectory service include the connected dire connector namespace, metaverse and client as seen in Figure 9.

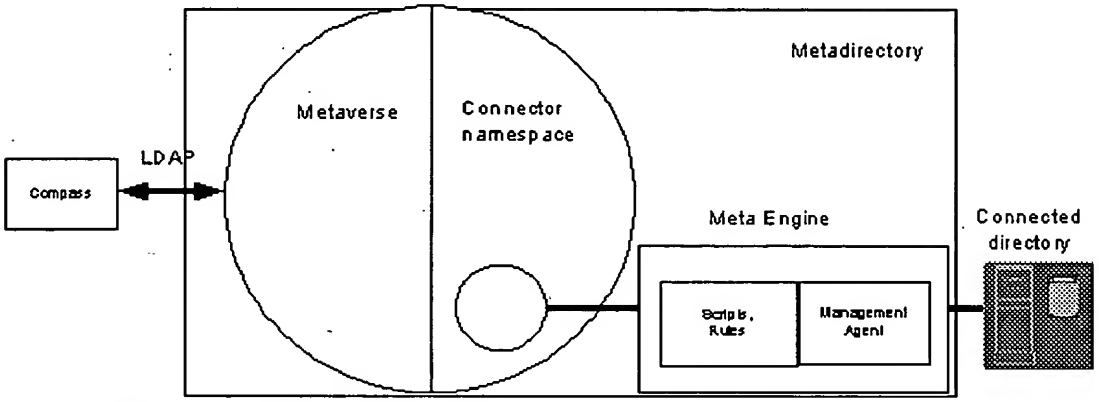

Figure 9 Microsoft Metadirectory Services Components

In this illustration, the Compass client is the administrative user interface that speaks LDAP to the metad supports the HTTP protocol for convenient end-user access through Web browsers.

The Metadirectory Namespaces

The metadirectory is broken into two namespaces.

- The *connector space* is the area into which connected directory entries are first imported. Each cor its own area in connector space. Connector space is a collection of special objects called *connector* The difference between these two object classes is that a connector has an attribute filled in with th Name of the metaverse object it is connected to. A disconnector does not have this attribute filled i disconnector exists in the connector space merely as a placeholder to represent an entry in the con corresponding metaverse entry may or may not exist. The connectors establish a link between an c metaverse and one in a connected directory, allowing synchronization and attribute flow. Disconn synchronization. Connector space objects always appear under Management Agents in the director usually have many attributes populated. They primarily function as an intermediary between the m particular connected directory.
- The *metaverse* is that portion of the directory that presents the integrated view of joined objects fr connected directories. Most metaverse content comes from connected directories. But it is also po metaverse objects with no connection to any connector space object or connected directory.

Consider the namespace represented by the following diagram:

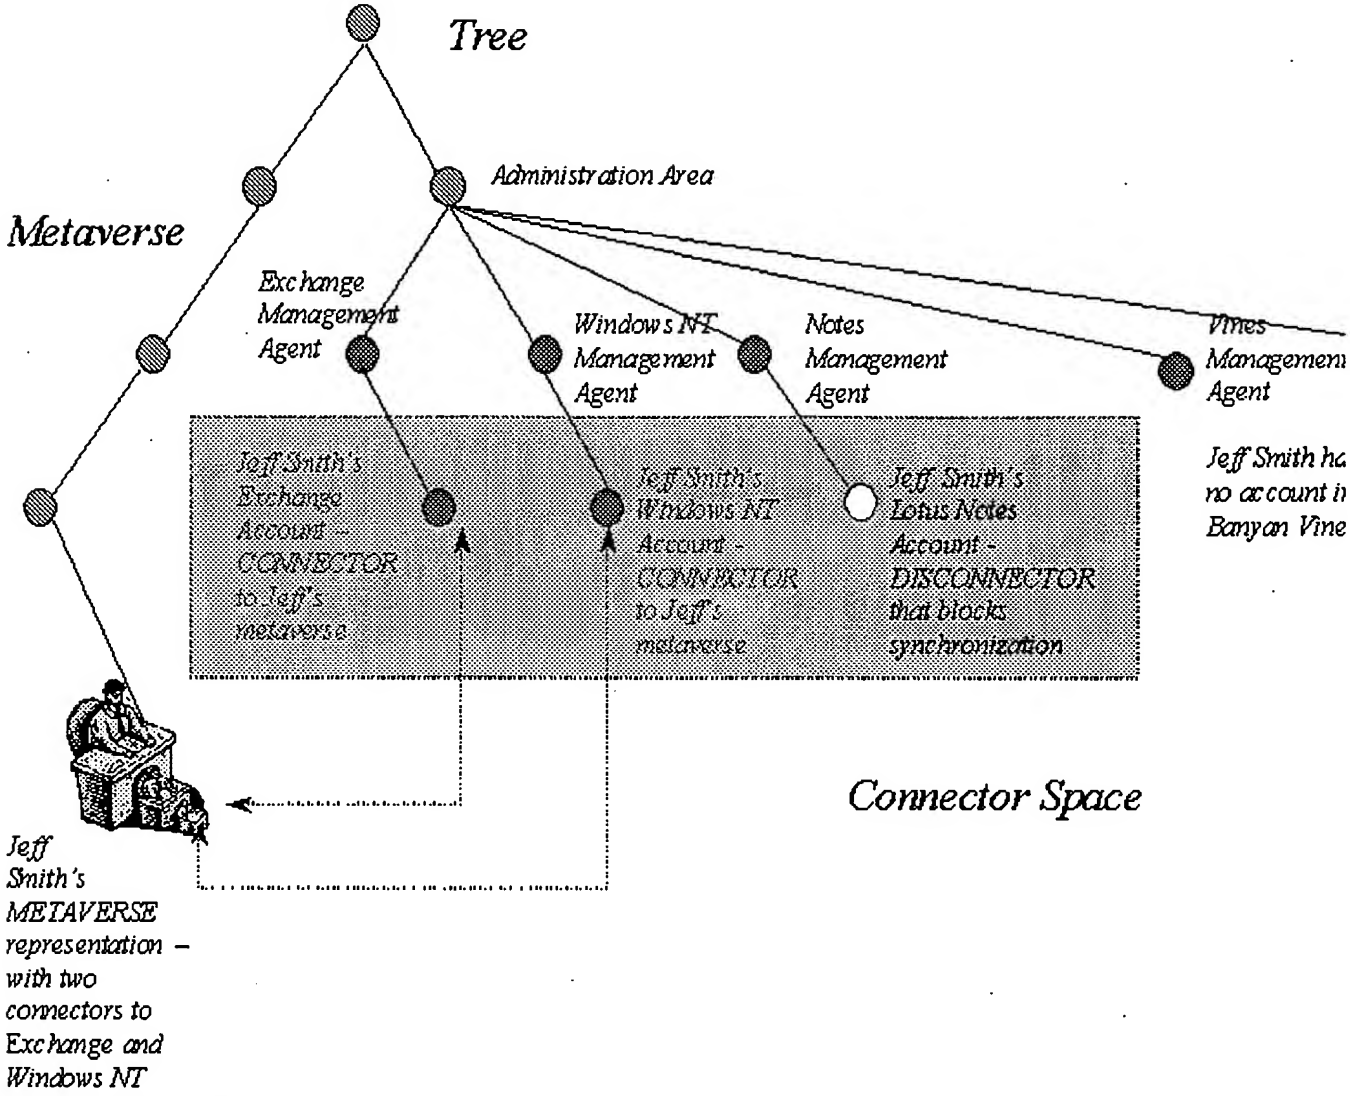

Figure 10 The Namespace

In Figure 10, the object representing Jeff Smith in the metaverse contains properties from the objects rep in Microsoft Exchange and the Windows NT? operating system. At one point, the object representing Je was joined to the metaverse representation; this object has since been disconnected. There is no Jeff Smi Also, Jeff Smith does not participate in the hire/fire scenario that is possible through TAMA (the Togeth Management Agent). TAMA is discussed in more detail later in this document.

In Figure 11 below, the Compass screen shots show the two namespaces side by side. The first is a view metaverse. It begins at the top ?The Known Universe ? and shows several branches of the metaverse tree Daniel Penn. The last entry visible on the screenshot is mdserver, which represents the metadirectory ser screenshot shows the hierarchy beneath the mdserver entry in more detail. This is where Management A connector space are located. It is also where the schema is defined for the metadirectory, where replicati information is kept, and where the default Administrator entry that controls the entire metadirectory is cr

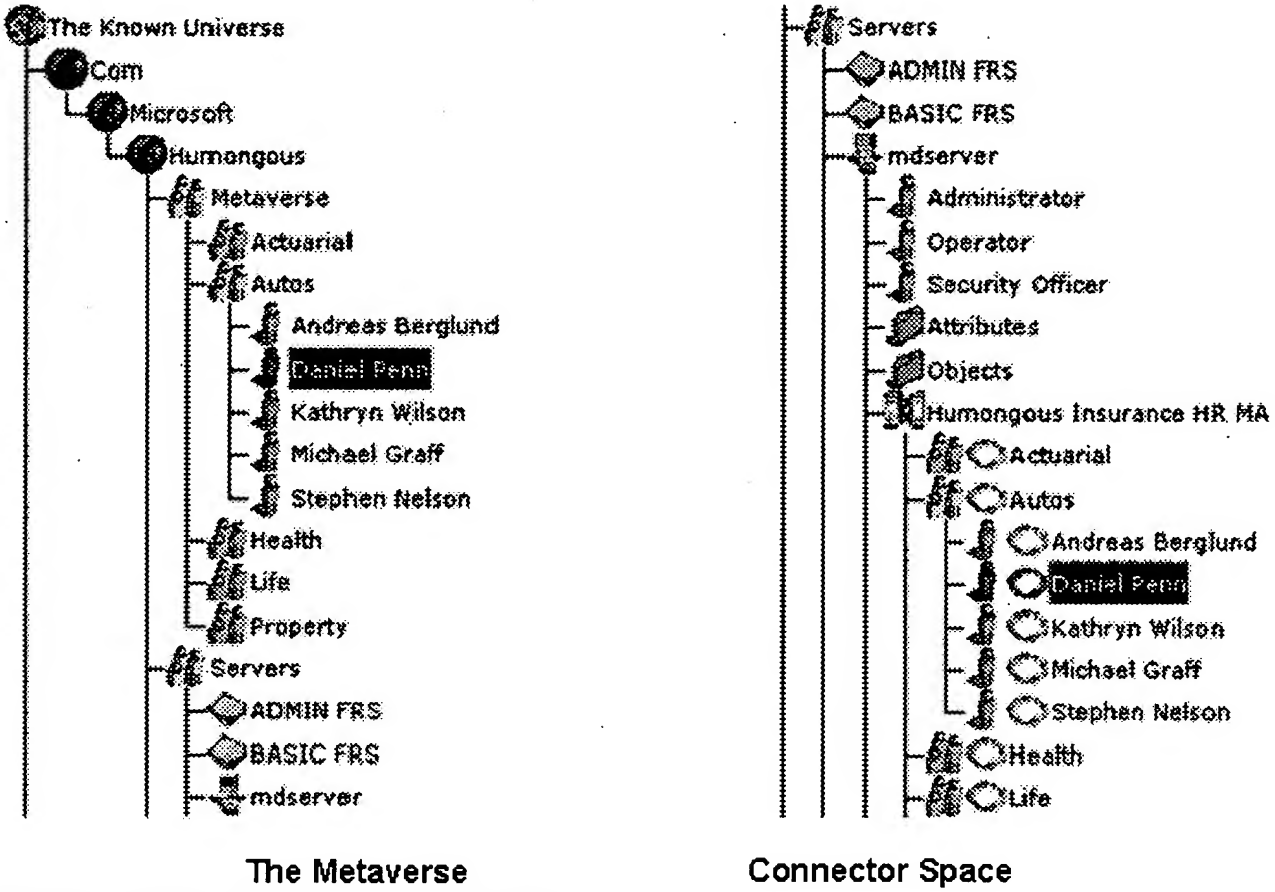

Figure 11 Screen Shots of the Namespace

In the illustration, objects called Autos (a department) or Daniel Penn (a person) exist in both the connec metaverse. Those in the connector space are connectors, and they are distinguished from the correspondi metaverse object by a special icon. The connector space for this connected directory is that area under Tl Insurance HR MA; other MAs, for the e-mail system, for example, might also contain a connector for Di connector space. The metaverse Daniel Penn object could hold attribute information from all of these sou

Management Agents

The Meta Engine controls the interactions between a connected directory and the metadirectory. It conta required to handle object creation and deletion, property integrity and history. It resolves property owner oscillation. These Meta Engine instructions are embodied in the metadirectory as the Management Agen specialized objects containing the configuration parameters, control scripts, transformation rules, attribut rules that define how a connected directory will be integrated with the metadirectory.

The MAs manage the relationships between connected directories and the metadirectory's connector nam metaverse at both the object and attribute level. They reside on the MMS server and are connected direct the internal configuration of the MA is different for each connected directory. An important note is that I require you to install additional software on any of your connected directories or other systems.

The Synchronization Cycle

The MA is a directory object and service that sets up directory synchronization. It defines how the synch

performed, and it performs the synchronization. A control script directs three separate phases of MA ope synchronization and update phases. These phases are illustrated in Figure 12 below.

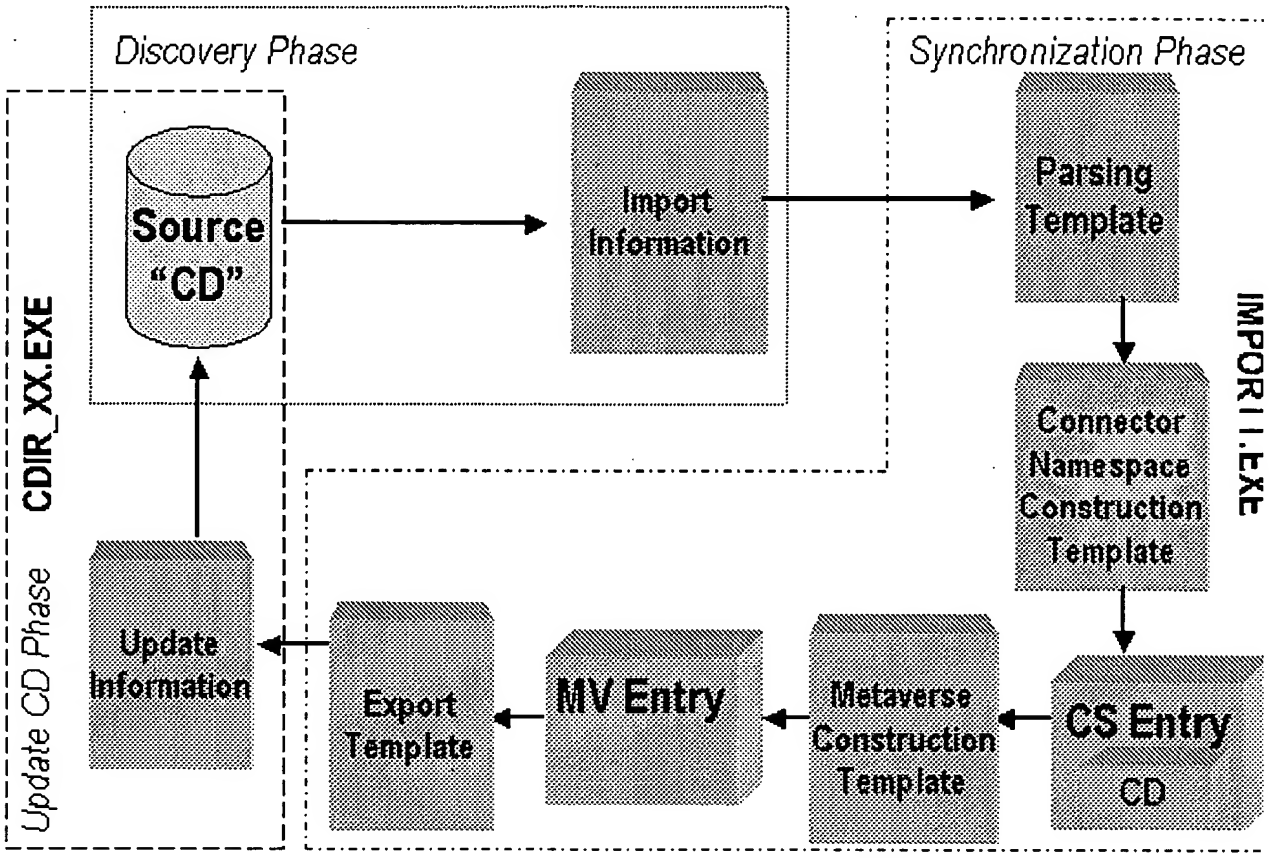

Figure 12 Management Agent Synchronization Phases

The discovery and update phases typically share code to bind to a connected directory and read and write information. The synchronization phase is the real heart of the MA. The synchronization phase uses imp templates and attribute flow rules (which are stored as properties of the MA object) to determine the exte changes that must be applied to both the connected directory and the metadirectory.

MMS comes with MAs for the most commonly encountered types of identity information repositories su directory, Windows NT, Microsoft Exchange, Banyan VINES, Netscape's directory service, Novell NDS Notes and cc:Mail to name but a few. Optimized support for Active Directory has also been added. By d MAs that come with MMS handle most of the common attributes that pertain to a given vendor's directo MA, for instance, maps the Notes OfficeTelephoneNumber attribute to the LDAP telephoneNumber attri Notes organizational structure to build a default hierarchy, and so on.

By modifying the scripts and templates, you can easily customize the supplied MAs to reflect any minor implementation of a connected directory type. For example, Exchange sites often use Exchange Custom specific information not included in the default Exchange directory schema. It is quite easy to customize Exchange sees Custom-Attribute-1 as specialTelephoneNumber in the metadirectory, and manages it acc

New MA types can be written using the information in the Management Agent Toolkit manual. New sou information, many of them outside the traditional network operating system (NOS) and e-mail directory

integrated and managed by MMS. As long as the repository ? a database, for instance ? can export inforr can easily create an MA to read that file and synchronize the repository and the metadirectory.

Managing Changing Information

When identity information about a person (or other object) exists in one or more connected directories ar in the metadirectory, who maintains it? If changes are made in both the metadirectory and the connected will soon drift out of synchronization. MMS allows you to determine not only where objects can be creai also in which directory individual attributes of existing objects can be modified.

MAs are scheduled to periodically compare the contents of the connected directory with the contents of t the contents differ, the MA synchronizes them. The connected directory and the metadirectory can differ

- Objects may exist in one that do not exist in the other.
- Objects that exist in both may have different attribute values.

The MA reconciles these differences and keeps the two directories synchronized according to the config synchronization rules you establish.

Managing Objects

The MA *operating mode* determines where the creation and deletion of a metadirectory object is manage connected directory (local management) or at the metadirectory (central management). As illustrated in I operating modes can be:

- Reflector. Additions and deletions in the connected directory are reflected in the namespace and m
- Creator. Additions and deletions in the metaverse are automatically performed in the connected di
- Association. Additions and deletions in the connected directory appear in the namespace but are n metaverse.

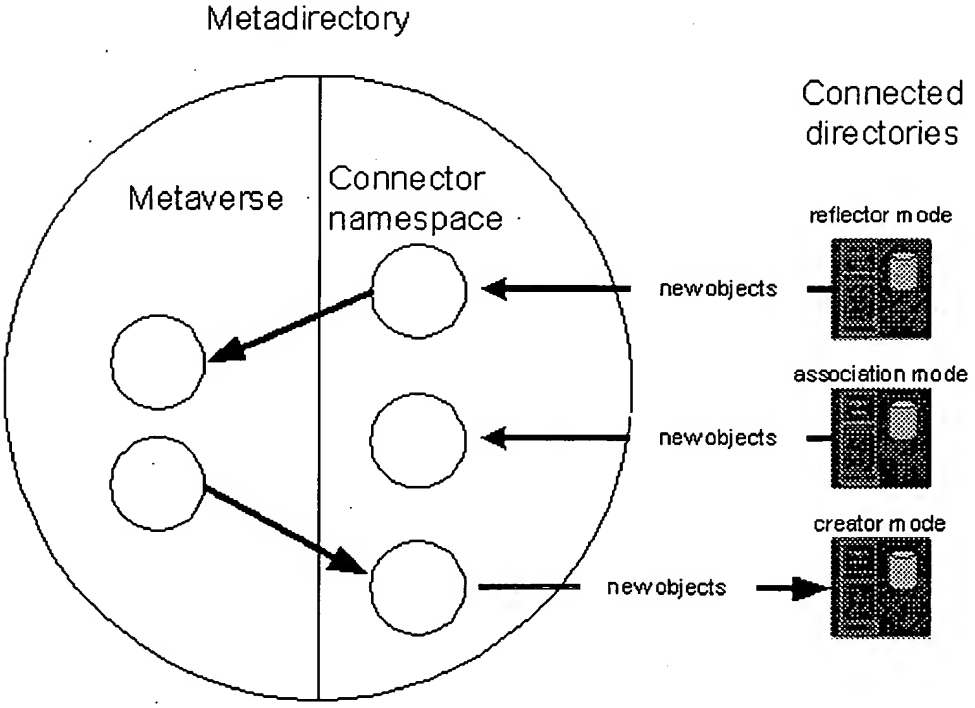

Figure 13 Management Agent Operating Modes

Local Management

When the MA is operating in *reflector mode*, the metadirectory simply reflects additions and deletions m directory. *Association mode* is a special form of reflector mode in which the two directories are associate Connected directory identity information is contained in the connector namespace, but it is not merged w information in the metaverse. Association mode is generally a transitional step to a reflector or creator m review the imported data before trying to join the connected directory objects into the metaverse.

Central Management

When the MA is operating in *creator mode*, objects can be created or deleted only in the metadirectory. deletions are then automatically performed in the connected directory. Should the connected directory ge metadirectory, the MA will automatically re-synchronize by adding or deleting connected directory obje

Managing Attributes

The MA operating mode determines only where *objects* can be created or deleted. *Attributes of existing* modified in either the metadirectory or the connected directory, regardless of the MA operating mode. W differ, an *attribute flow rule* specifies whether the metadirectory or the connected directory is authoritati attribute flow rules can take effect, the connected directory object, through its connector space entry, mu metaverse entry.

The Join

The join establishes a link between a metaverse object and a specific connector space object. In linking t to the connector space object, the join indirectly also links it to the connected directory. Objects can be j according to predefined join criteria, or interactively by the administrator.

A variety of join options is necessary for two reasons. First, there may be connected directories with entr merged in a common metaverse object. Second, there may be no sure way of knowing when a metaverse particular connector space entry describe the same object. For example, there may be several Jeff Smiths each represented in a different connected directory. One person may appear in those directories under se Jeff Smith, J. Smith, Jsmith or Smith, Jeff Q. Some degree of administrator intervention is often required kinds of ambiguities. But it is also true that in many cases there is no ambiguity. You can simply match name or other attribute (employee number, for example).

There are, in fact, three different ways in which objects can be joined.

The *Join Action* in the Compass client lets you define an automated batch join, based on predetermined j execute it whenever you want. This batch join is normally used when you first bring a new connected dir metadirectory ? perhaps in Association mode. Inevitably you will be left with exceptions ? those connect that fall between the cracks of the join criteria. They are left un-joined and remain *disconnectors* rather t can then use the stand-alone *Account Joiner* application to deal with these exceptions, by searching the n or by creating a new metaverse object corresponding to the disconnector. But new objects may appear in directory at any time. These also may have to be joined to metaverse objects. For this kind of ongoing m configure the MA to automatically use join criteria when it creates a new disconnector (an un-joined con If there is no ambiguity, it can perform the join at the time of creation.

If you are going to allow joins to be made automatically, you must first define the rules ? the join criteria Action and the MA join offer you a Configure the Join interface where you can do exactly that. See Figu example.

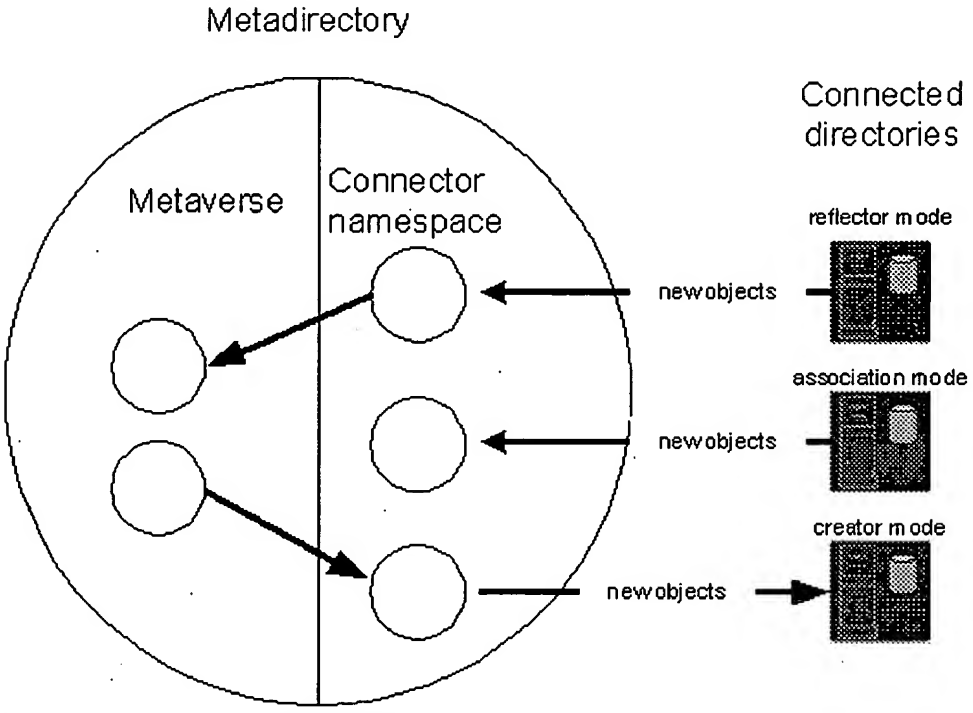

Figure 14 Configuring Join Rules

The Join Action looks at each disconnector in the connector space and searches for possible matches in t on search attributes you specify. This search may return several matches. The Join Action then applies a to determine which, if any, of the possible joins it should accept. If it finds a suitable match, it establishe two entries automatically. Choosing the option, *Try to join before reflecting new entries,* tells an MA in i the same join criteria (in addition to its usual techniques) to search for an existing matching metaverse ol (reflecting) a new one. The Account Joiner, on the other hand, lets you define rules as you go, experimer search criteria until you find a matching object or decide to create one. You may then want to incorporat search techniques into your join rules to handle such cases automatically in the future.

Attribute Flow Rules

When multiple MAs update the same metaverse object, their attribute flow rules must define which conn controls each attribute. If not, they will overwrite each other's changes. You must define which connecte authoritative source for each attribute.

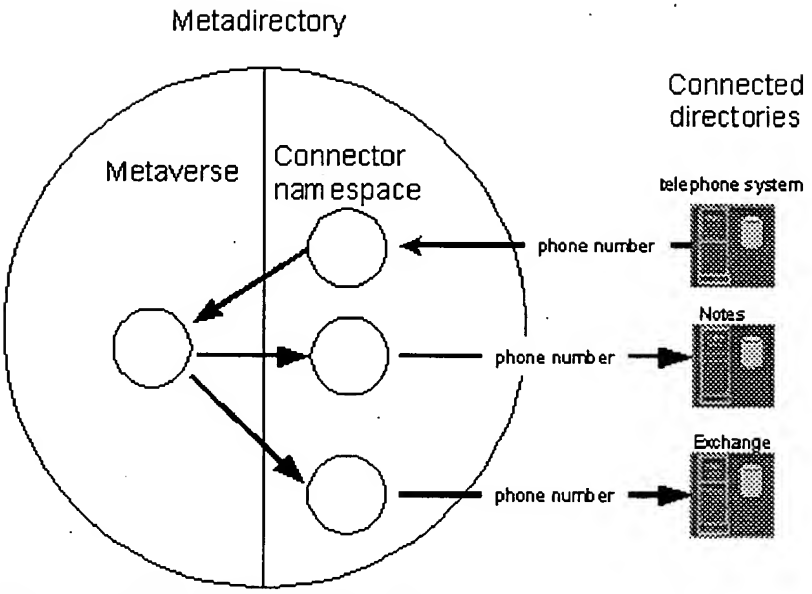

Figure 15 Metadirectory Attribute Flow

In the illustration in Figure 15 above, the telephone system is the authoritative source for a person's phon MMS is set up so the telephone system connected directory is the only one that can modify the phone nu to the phone number from the telephone system will be synchronized to other connected directories, suc Notes. If a user or administrator tries to change the telephone number in these connected directories, the the telephone system will overwrite the changes.

These attribute flow assertions can be made within a simple point-and-click interface. See Figure 16 belc

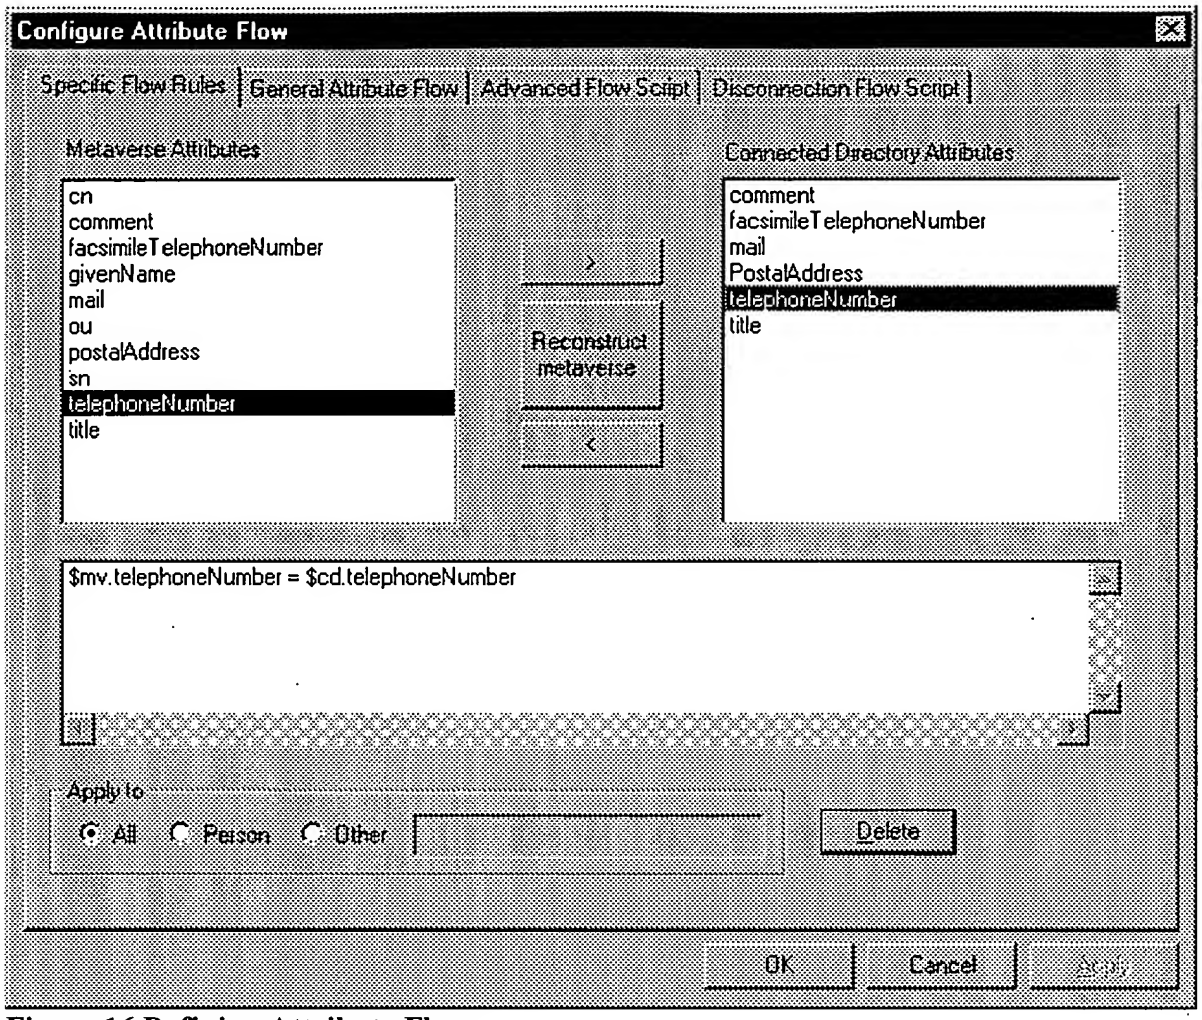

Figure 16 Defining Attribute Flow

By simply selecting the attribute involved and clicking a directional flow button, you establish a simple r system MA:

$mv.telephoneNumber = $cd.telephoneNumber

This assertion states the phone number attribute in the metaverse should be set to the phone number attril telephone system's connected directory.

In the Exchange and Notes MAs, the corresponding attribute flow rule would be something like:

$cd.telephoneNumber = $mv.telephoneNumber

This assertion states the phone number attribute in the connected directory should be set to the value of t attribute in the metaverse.

An Advanced Flow Script lets you deal with more complicated flow rules using simple script like condit

This level of control over attribute flow greatly enhances the metadirectory's distributed management ca identity information to be maintained where it makes the most sense. Global information (for example, e can be maintained centrally. And local information (for example, phone numbers) can be maintained loc

Data transformations

MAs import directory update files from their connected directories and send update files back to them. C simple of MAs is direct attribute transfer sufficient. Among the complicating factors are:

- There may be no exact metadirectory equivalent of certain connected directory attributes.
- The connected directory information may be in a different format than that required by the metadi
- It may take more than one attribute within the connected directory to match a metadirectory attribu
- Additional metadirectory attribute information will almost certainly be necessary to ensure that the within the whole metadirectory.
- The connected directory may contain objects that you do not want to import into the metadirectory

MAs use templates to determine how to input and output attribute values. Templates are written in a higl language that the *importt* program interprets and acts upon when importing or exporting metadirectory d

The MA template language therefore provides the capability to:

- Perform simple direct modifications on attributes.
- Use built-in functions to transform attributes.
- Obtain additional information from other objects in the metadirectory.
- Provide control over template execution through conditional control structures.
- Define metadirectory objects to be included or excluded in a directory update.

The leftmost portion of **Table 1** below shows a record exported by the cc:Mail Export/Import utility duri phase for import into the metadirectory. On the right is the parsing template that describes it in terms of attributes and temporary variables. You can see in general terms how attribute substitution is defined.

Table 1

File Contents	Template
Name: Dunn, Matt	Name: $v_surname,$v_givenName
Locn: L	Locn: $cd.zcCcLocation
Addr: ccmPO	Addr: $cd.zcCcPostOffice
Cmts:	Cmts: $cd.description

It is evident that this connected directory does not export very much attribute information. We need far n construct a full metadirectory entry. We need values for all the attributes that make up the entry's Disting example, as well as its object class. These other attributes must be constructed from identity information and other information known to the MA. Each MA, therefore, also has a set of construction templates to parsing templates.

The following excerpt from a construction template suggests how this kind of information is created.

Table 2

```
If $cd.zcCcLocation = P (i.e., it is a PostOffice entry)
then
$mv.zcoc = organizationalUnit (object class)
$cs.zcoc = zcCcMailPostOffice,zcAliasThing,Top
$mv.organizationalUnitName = $v_surname(,$v_givenName)
. . .
```

```
else
$mv.zcoc = zcPerson
$cs.zcoc = zcCcMailBox,zcAliasThing,Top
$mv.commonName = $get_substring("$v_givenName\
(^ $v_surname)", "", "@"
. . .
endIf
```

You do not have to understand the details of the template in **Table 2** to grasp how it can be used to contr information gets in and out of the metadirectory, some directly from the connected directory, some from

It is clear that a reflector MA creating a new metaverse object would use the construction template to do corresponding metaverse object already exists and is joined to the connector? What about the attribute fl is that the construction templates define projected or potential attribute values which might have to be se rules determine which ones actually are set by a particular MA. The flow rules thus supplement the temp the construction template in our example assigns a value to the common name attribute of the metaverse ($mv.commonName) based on information from the connected directory. It must to reflect a new metave situation that the attribute flow rules come into play. They are not separate from the templates but supple

A metaverse object, however, is not simply a collection of all the attributes of all its connected directorie you present to the world only the information the world needs to see. Easily modifiable templates and flc be selective and discriminating. You create metaverse objects, depending on the use of the metadirectory directories remain in place, performing their original roles. The metadirectory goes beyond the connecte not replace them ? unless you choose to.

Enabling the Hire/Fire Scenario

A key role for a metadirectory service with enterprise customers is programmatically updating access to resources during an employee's term with a company. When an employee is hired, he or she will require resources such as files and printers. The employee also needs services such as e-mail. When the employe goes through reorganization, different access rights might be needed and different services might be requ

MMS supports this scenario through its Together Administration Management Agent (TAMA). TAMA i MA that can manage and coordinate the activities of several other standard MAs. This capability allows multiple connected namespaces to ensure that connectors, and by extension, accounts are created in the c Typically its activity is initiated when new entries are created in a particular connector namespace. TAM corresponding connectors under different MAs to provision accounts in different connected directories. I entry in the HR connector space could cause a new Windows NT account to be created, an Exchange ma and perhaps a Lotus Notes ID to be created as well. Using the same scripting language as other MAs, TA precise control over when and where these new accounts will be created. Conversely, if someone leaves is, they are removed from the HR system), TAMA can ensure that all of the associated accounts in the di directories are cleaned up and deleted.

Putting Metadirectory Services to Work

The following examples show how MMS has been put to work in the real world to solve real enterprise p out how the metadirectory has been used to implement the two key scenarios enabled by MMS: change c control. The example companies, organizations, products, people and events depicted herein are fictitiou with any real company, organization, product, person or event is intended or should be inferred.

Northwind Traders

Northwind Traders is a conglomerate of companies, large and small in a worldwide holding corporation. use different kinds of systems, including Lotus Notes, Netscape, GroupWise and Exchange. In their met these systems are connected to the head office metadirectory server both through a WAN and over the In

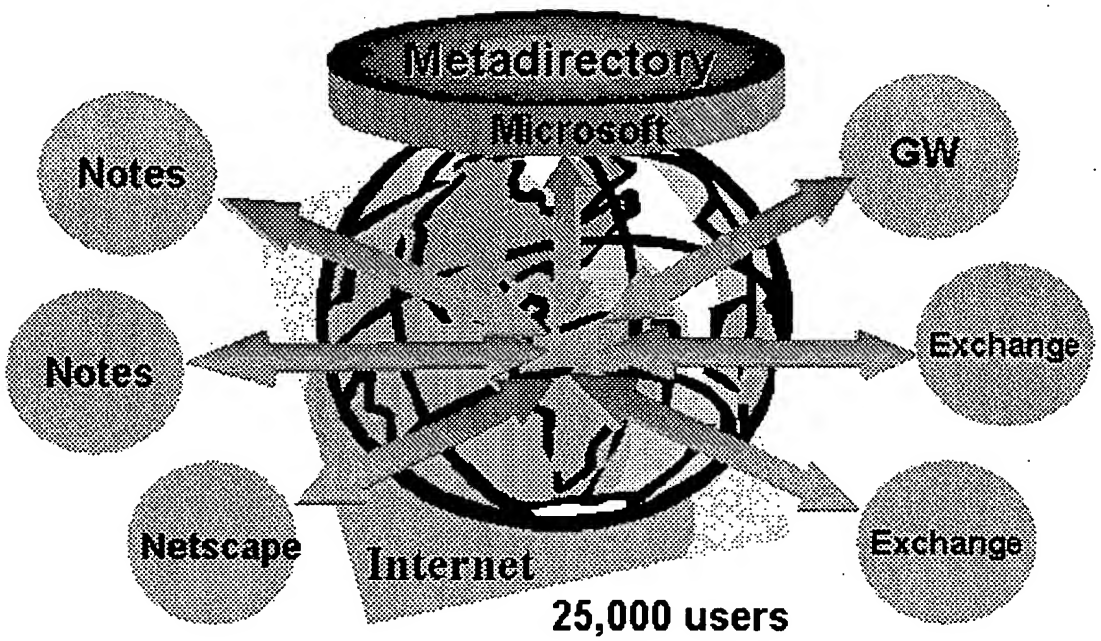

Figure 17 A Conceptual View of Northwind Traders

Figure 17 above provides a conceptual view. However, the simple illustration above does not portray the political implications involved in managing such an environment.

In England, there is a metadirectory service which brings together a number of different systems that are companies owned by Northwind Traders. And because it is a holding corporation, there are many, highly business units, not only in England but around the world. The people who run the metadirectory service able to bring the different systems together and even integrate them with the central HR system. Similarl Egypt, and in a number of other countries, the affiliated companies run different e-mail packages like Ex whose attribute flows are integrated with the metadirectory in England.

Connected directories are managed locally and reflected in the metadirectory. The metadirectory service directory synchronization engine enabling local users to see the entire corporate address book. Only the actually access the metaverse directly.

With Microsoft metadirectory technology, Northwind Traders did not have to deploy expertise in all of t all of these countries. It was able to set up the connections over the Internet so that, without any local ex Traders could achieve the desired flow of information and attributes.

Interestingly, the holding company also holds an American corporation which is on virtually the same sc corporation. The Americans had no intent of sacrificing their local autonomy to the British corporation s metadirectory service was installed in the U.S.

The following diagram, Figure 18, better represents the current reality within Northwind Traders.

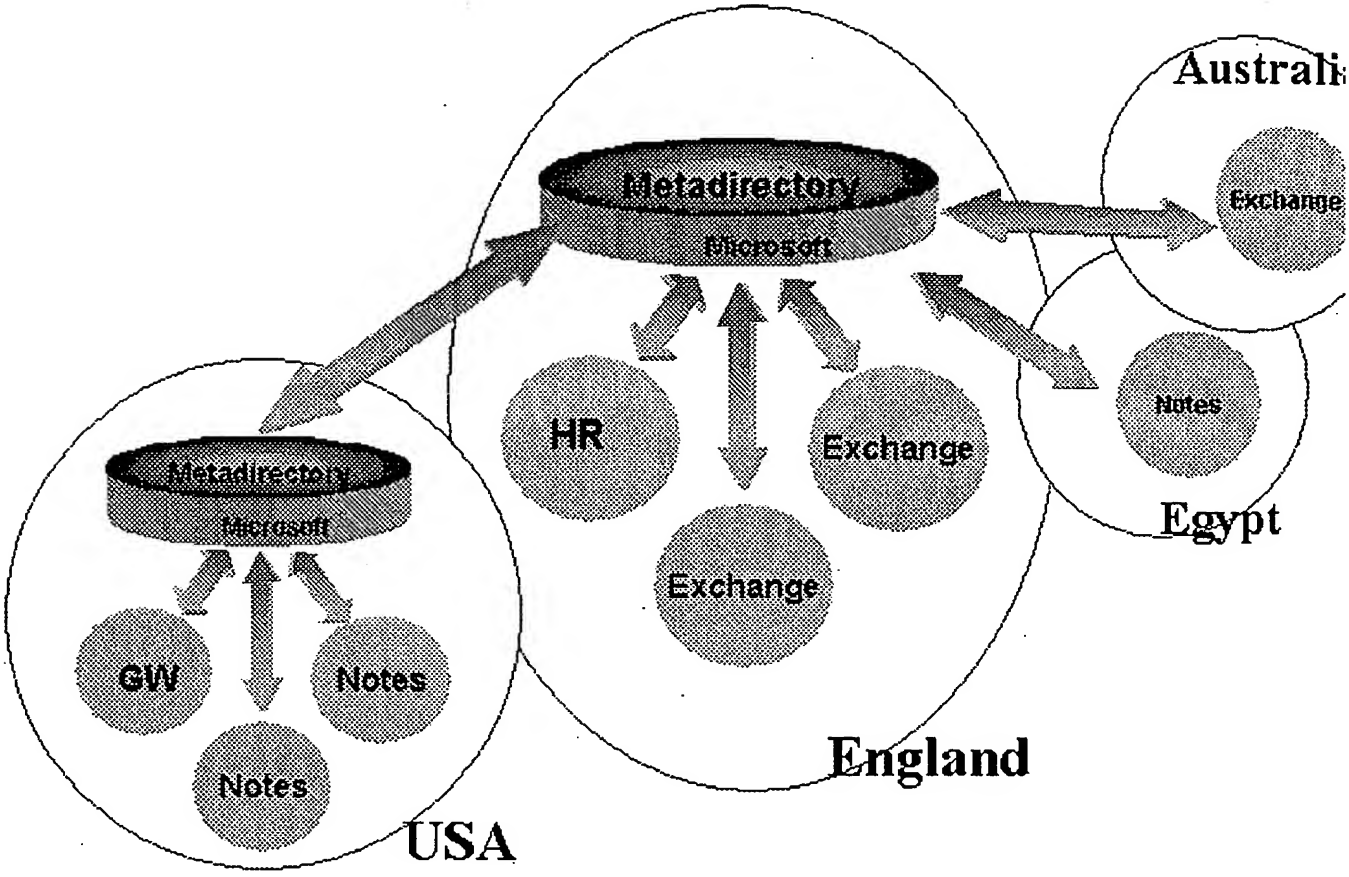

Figure 18 A More Realistic View of Northwind Traders

Not only is there local management at the connected directory level, but there is also regional manageme metadirectory server level. The two servers each manage part of the world, and then they exchange infor metadirectory replication. So each contains a complete representation of the corporate directory. The resi whether in Cairo or in Chicago, is a full, accurate and always up-to-date address book. And the administi metaverse, a seamless, integrated view of the entire company in all its complexity.

Coast Appliances Corporation

The Coast Appliances Corporation is an organization that required a hire/fire scenario in which HR coul people in the metadirectory. The metadirectory would then integrate the information on each person whc within the telephone system, the e-mail system, RACF, certificate systems, and other data stores through While conceptually simple, this project became very complicated because it had to deal with the messine

When the Coast Appliances Corporation began this project, there was no way to easily create a join betw information that resided in each of these data sources. The data was very difficult and very dirty. In fact, interesting things about the person information was that only 65 percent or less was found in any of the s is even more interesting because the person information wasn't the same 65 percent in any of the systems percent was in telephone, another 65 percent was in HR (because there were a lot of contractors and peo systems). And 40 percent was in RACF.

How does the Coast Appliances Corporation perform a join in a situation like this? If a rule is set up to u attribute from the HR store, then the people from the telephone system who aren't in HR aren't joined. Tl

the people within the telephone system.

The Coast Appliances Corporation needed a fairly sophisticated way to join the people information. The with the HR system and brought its 42,000 employee objects into the metadirectory to provide the base. Appliances Corporation focused on the telephone system which had 45,500 user objects. When the telep objects were imported into the metadirectory, 34,000 of them were automatically batch-joined with exist in the HR system. But 11,500 of them were not represented in the HR system and, consequently, the met the HR MA and the telephone MA were running in reflector mode, 11,500 new people objects were add The telephone MA was configured to *Try to join before reflecting* and the join rule was to join on comm telephone system objects whose common names matched names already in the metaverse (from the HR to them. For the rest, new metaverse objects were reflected using the telephone system common names. however, simply accepted and merged. Those new objects reflected by the telephone MA were initially part of the metadirectory tree where they could be examined, accepted or rejected, and eventually moved organizational location.

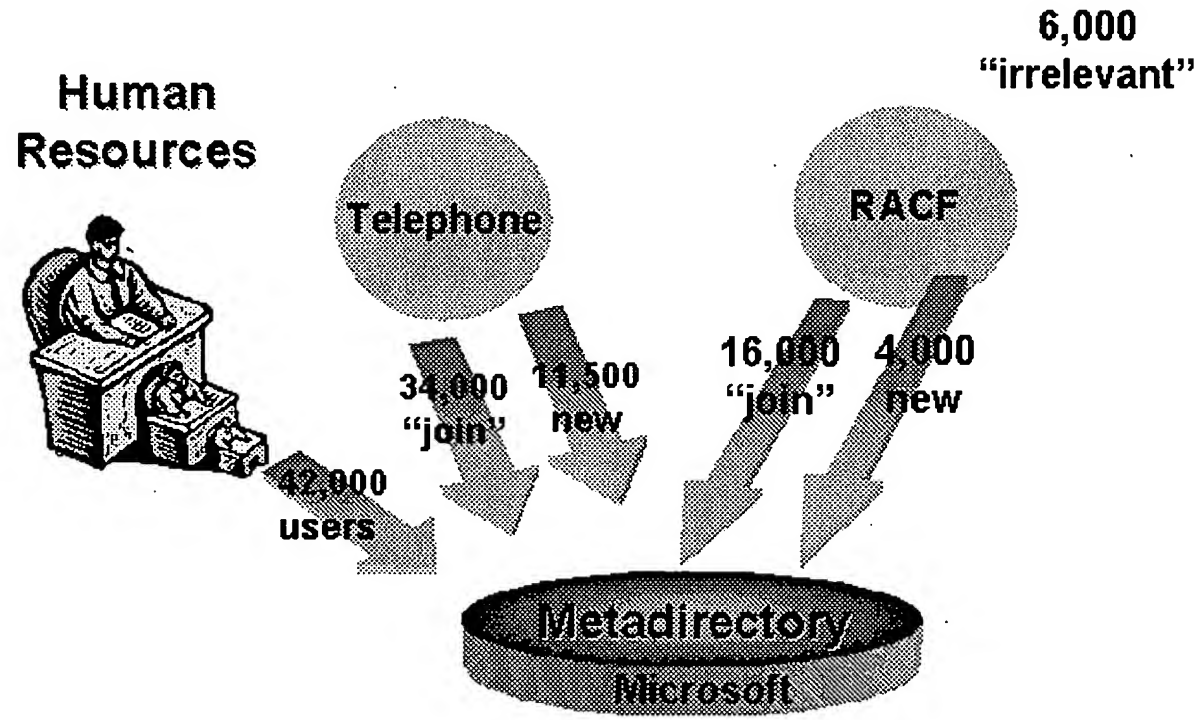

Figure 19 Coast Appliances Corporation's Metadirectory Join

After integrating the HR and telephone systems, the Coast Appliances Corporation turned its attention to Based on the join rules noted above, 16,000 of the user objects were batch-joined. Eventually 4,000 beca metaverse. Some entries represented persons, others represented functions or roles. It was interesting tha objects in the RACF account system were found to be irrelevant. They belonged to users who no longer with the Coast Appliances Corporation. Many of the jobs that were actually running in the computer cen run under the permissions of users who had left the organization. And a certain amount of mainframe ca reclaimed because of the rationalization process which explored all of these irrelevant and unused accou illustrated in Figure 19 above.

By the end of identity information integration with just these three systems, the Coast Appliances Corpo rationalize 49,500 identity objects. It then tackled the other systems and progressively made sense out of of this process, the Coast Appliances Corporation had integrated its identity information and found a trer

erroneous information in all of these systems. Also, it discovered that some of the identity information w control of the people who managed it nor under the control of the people who were being referenced.

Obviously, joining identity information across isolated data sources within a large enterprise is not alway metadirectory services can automate much of the process and highlight problems and inconsistencies. Th installing a metadirectory solution brings its own benefits and cost savings because it allows you to clear dirty information. Once fully implemented, the metadirectory service makes many of those savings ongo reliable, integrated identity base on which other initiatives and applications can be built. For example, th Corporation took advantage of the metadirectory to centrally generate unique IDs for every employee an of the local administrators, flow that attribute back to all the connected directories. Such a program woul unmanageable if not unthinkable without the metadirectory.

Metadirectory Services Enhances Active Directory

Microsoft's goal is to use metadirectory services to enhance Active Directory, providing a comprehensiv platform. Many Windows 2000-based applications are Active Directory enabled. End-users can find the the location-based printing feature. Applications and service policies can be centrally managed from Act then downloaded to a group of end users. Exchange 2000 replaces the Exchange directory with Active D DNS (Domain Name System) is tightly integrated with Active Directory. These are a few of the great int Microsoft's enterprise customers can enjoy with Active Directory deployments.

Microsoft Metadirectory Services enhances Active Directory by providing such services as:

- Synchronization of multiple connected directories within a centrally managed hub-and-spoke mod
- Programmatically joining multiple views of an object into one unified view. Although it is unreali enterprise's identity information can be programmatically joined 100 percent of the time, MMS's f view of identity information greatly simplifies this important step.
- Setting a connected directory as the authoritative source for an attribute.
- Integration with the business process through support of the hire/fire scenario.
- A simple and flexible environment that allows short scripts to be added for customization within a environment.

Future releases of MMS will further integrate with Active Directory while at the same time being enhanc customer scenarios. The MMS 2.2 release (available July, 2000) provides an optimized Active Directory that takes advantage of the Active Directory advanced replication protocol to detect changes and copy th real-time directly into Active Directory. Hence, Active Directory can be used as the primary administrati metaverse objects.

Future integration plans with Active Directory include integration with the Windows 2000 authentication overall integration of MMS within the Windows 2000 Server platform.

The combination of Active Directory and MMS is a compelling solution for a distributed systems platfo accounts. Microsoft Metadirectory Services is a great enhancement to Microsoft's Windows Server distri offering.

Summary

Managing identity data in a modern enterprise network presents many challenges. Identity data comes in may be scattered in several repositories. A metadirectory collects all the identity data in one place and pr

managing the data, regardless of its format. A metadirectory allows a business to reduce the cost of admi duplication of data, reclaim wasted network capacity, resolve discrepancies in the data, and make the ide conveniently available. Microsoft Metadirectory Services, which evolved from ZOOMIT VIA, provides solution for the identity management challenges faced by modern enterprises. It allows a geographically to manage data at the local level, or at a central location. Or, an enterprise may choose a combination of management. The metadirectory can provide an accurate, always up-to-date record of information about things as addresses, phone numbers, e-mail, departmental titles and document files.

For the latest information on Windows 2000 Server and Active Directory, check out our Web site at http://www.microsoft.com/windows2000/technologies/directory/ad/default.asp.

For further information on Microsoft Metadirectory Services visit http://www.microsoft.com/windows2000/technologies/directory/default.asp.